arXiv:1511.08049v1 [cs.SE] 25 Nov 2015

Industrial Experiences with a Formal DSL Semantics to Check Correctness of DSL Transformations

Sarmen Keshishzadeh[1], Arjan J. Mooij[2], and Jozef Hooman[2,3]

[1]Eindhoven University of Technology, Eindhoven, The Netherlands
[2]Embedded Systems Innovation by TNO, Eindhoven, The Netherlands
[3]Radboud University Nijmegen, Nijmegen, The Netherlands

Abstract

A domain specific language (DSL) abstracts from implementation details and is aligned with the way domain experts reason about a software component. The development of DSLs is usually centered around a grammar and transformations that generate implementation code or analysis models. The semantics of the language is often defined implicitly and in terms of a transformation to implementation code. In the presence of multiple transformations from the DSL, the consistency of the generated artifacts with respect to the semantics of the DSL is a relevant issue. We show that a formal semantics is essential for checking the consistency between the generated artifacts. We exploit the formal semantics in an industrial project and use formal techniques based on equivalence checking and model-based testing for consistency checking. We report about our experience with this approach in an industrial development project.

1 Introduction

A domain specific language (DSL) [28] abstracts from implementation details and is aligned with the way domain experts reason about a software component. By focusing on the essential concepts in a problem domain, DSLs facilitate the involvement of domain experts in the development of DSL specifications.

Tool support for the development of DSLs is improving constantly. Language workbenches such as Xtext [4] enable language designers to define their languages and develop transformations that generate various artifacts from domain specific models. This has boosted the popularity of DSL approaches in industry; see [17, 29].

The development of DSLs is usually centered around a grammar and transformations that generate implementation code or analysis models from DSL specifications. In such a setting, the main focus is on the transformation to implementation code which is very valuable in industrial practice. The generated analysis models allow to reason about the behavior described in DSL models. For example, properties can be verified against verification models or simulation models can be used to explore the modeled behavior interactively.

In DSL approaches, the semantics of the language is often defined implicitly and in terms of the generated implementation. Although DSLs focus on the essential concepts of their respective domains, the semantics of the language constructs are not always obvious. The lack of a formal semantics can give rise to different interpretations and this can potentially cause inconsistencies between the transformations.

Various authors [6, 9] have proposed to use a formal semantics to describe the precise meaning of the language constructs. The formalization also allows them to have a single reference that should be followed by all the transformations.

Implementing a transformation from a DSL to a programming/modeling language is not a trivial task. Having a formal semantics as a reference does not guarantee consistency between multiple transformations from a DSL. The developer of a transformation should have a deep understanding of the DSL and the target language and construct a transformation that does not deviate from the semantics of the DSL. In [27] the authors indicate that such tasks are very error prone and introducing redundancy is an effective way to reduce the rate of mistakes.

We propose to use the formal semantics of a DSL and introduce redundancy using the following formal techniques:

- equivalence checking (by transforming models to a formalism that allows checking equivalence of behaviors);
- model-based testing (by testing conformance of executable models to a test model).

In this paper, we report on our experiences with using these redundant mechanisms in the context of an industrial DSL. The results obtained in the project show that these formal techniques can effectively detect inconsistencies between transformations from a DSL. To hide the complexities of these techniques, we have developed a push-button technology that allows industrial users to automatically perform these checks for the artifacts generated from a specification in the DSL.

Some authors [10] propose to formally prove the correctness of transformations. Since in practice transformations are improved regularly, we use automated techniques based on equivalence checking and model-based testing to validate the consistency between the artifacts, as opposed to formally validating the transformations themselves.

Overview We discuss preliminaries in Section 2. In Section 3 we describe an industrial control component and informally introduce a DSL for describing its behavior; the abstract syntax and semantics of the language are discussed in Section 4. In Section 5 we consider the formalized semantics and construct a transformation from DSL models to models that enable simulation. In Section 6 we generate test models from DSL specifications and apply model-based testing to assess the quality of an implementation of the industrial component. The correctness of the used models with respect to the semantics is validated in Section 7. In Section 8 we discuss about our experiences with different types of model transformations. In Section 9 we discuss related work. Section 10 contains conclusions and directions for future research.

2 Preliminaries

In this section we give an overview of the micro Common Representation Language 2 (mCRL2) [11]. mCRL2 is a process algebra that extends the Algebra of Communicating Processes (ACP) [8] with data and time. The mCRL2 language and its supporting toolset [2] can be used to specify and analyze the behavior of distributed systems and protocols.

In this short overview, we focus on the language constructs that we need throughout this paper. The interested reader can refer to [11] for a detailed explanation about this process algebra. We first explain the way data types are defined and used in mCRL2 (Section 2.1). Then, we describe behavioral specifications in the language (Section 2.2).

2.1 Data Specification

mCRL2 offers ways to specify data types (also known as sorts) and use their elements in specifications. Standard data types such as natural numbers ($\mathbb{N}$) and booleans ($\mathbb{B}$) are predefined in the language. Common operations on these data types are also available (e.g., $\approx$ denotes equality for $\mathbb{N}$ and $\mathbb{B}$).

Besides the standard data types, the user can define new data types in a specification. One way to declare a new data type is to explicitly characterize its elements in a structured data type. For instance, we can define a structured sort called *Color* with the elements *Red*, *Green*, and *Blue* as follows:

sort $\quad Color \;=\; \textbf{struct}\;\; Red \mid Green \mid Blue;$

It is also possible to declare structured types that depend on other data types. For instance, a data type called *Message* that contains pairs of natural numbers can be defined as follows:

sort $\quad Message \;=\; \textbf{struct}\;\; Pair(fst{:}\mathbb{N}, snd{:}\mathbb{N});$

Elements of *Message* have the shape $Pair(n_1, n_2)$ where $n_1, n_2 \in \mathbb{N}$. The declaration of *Message* provides two projection functions, namely, *fst* and *snd*. These

functions can extract the elements of $Pair(n_1, n_2)$:

$$fst(Pair(n_1, n_2)) = n_1 \qquad snd(Pair(n_1, n_2)) = n_2$$

Functions can also be declared and used in the mCRL2 language. Given two data types A and B, the notation $A \to B$ denotes the sort of functions from A to B. The mCRL2 language includes a specific operator called function update for unary functions. For a function $f \in A \to B$ this operation is denoted by $f[a \to b]$. The notation represents a function that maps a to b and maps all the other elements of A like f does.

We can declare the sort of functions from $\mathbb{N}$ to $\mathbb{N}$ as follows:

sort $NatFunc = \mathbb{N} \to \mathbb{N};$

We consider two examples of this sort:

- *succ*: given $n \in \mathbb{N}$ returns $n + 1$;
- *condsqr*: given $n \in \mathbb{N}$ returns n^2 when $n > 10$; otherwise, it returns n.

We declare these functions using the **map** keyword:

map $succ, condsqr : NatFunc;$

To define *succ* and *condsqr*, it is necessary to specify how calculations are performed in these functions. This is realized by introducing equations using the keyword **eqn**. Variables that are used in the equations are declared by the keyword **var**.

var $n : \mathbb{N};$
eqn $succ(n) = n + 1;$
$condsqr(n) = if(n > 10, n^2, n);$

The conditional operation has the shape $if(c, t, u)$. It evaluates to the term t if the condition c holds and it evaluates to the term u if c does not hold.

Lists are also a built-in data-type in mCRL2. The set of lists where all elements are from a sort A are represented by $List(A)$. Elements of $List(A)$ are built with two constructors: [] the empty list and $a \triangleright \ell$ which puts a (of type A) in front of list ℓ (of type $List(A)$). A list can be defined explicitly by specifying its elements and putting them between square brackets. For example, $[22, 4]$ is a list of natural numbers.

2.2 Process Specification

The mCRL2 language allows us to specify behavior using a small set of primitives and operators. We use a simple example to describe some basic constructs that can be used to specify behavior in mCRL2. The example is a modulo 3 counter that starts counting from 0 and resets itself when it reaches 3.

In mCRL2, behavior is described in terms of processes. Actions are elementary processes and they represent observable atomic events. Actions can carry

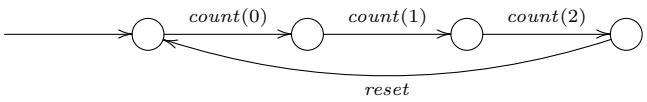

Figure 1: Behavior of *Counter*

data parameters. In our simple example, *count* and *reset* can be considered as actions performed by the counter. For example, *count* carries a data parameter to indicate the current number. These actions are declared as follows:

act $count : \mathbb{N}$;
$reset$;

Actions can be combined using different operators to form processes that specify more complex behavior. For instance, the non-deterministic choice between process p and q is denoted by $p + q$. The sequential composition of p and q is denoted by $p.q$. The resulting process first performs the behavior of p and then behaves as q. Data values can also influence the course of actions. Suppose c is a boolean expression. The process $c \to p \diamond q$ behaves as p if c holds and otherwise it behaves as q. The "else" part of the conditional operator can be omitted. If c does not hold in $c \to p$, deadlock will occur.

The following process expression specifies the modulo 3 counter. The process *Counter* carries a data parameter to keep track of the current number. The initial behavior is specified by **init**, i.e., counting starts from 0.

proc $Counter(n{:}\mathbb{N}) = (n < 3) \to count(n).Counter(n+1)$
$+ (n \approx 3) \to reset.Counter(0)$;
init $Counter(0)$;

For any $n \leq 3$ exactly one of the conditions $n < 3$ and $n \approx 3$ will evaluate to *true*. The process performs action $count(n)$ when $n < 3$ and then behaves as $Counter(n+1)$. When $n \approx 3$, the process performs action *reset* and starts counting from 0. Fig. 1 depicts the labeled transition system of the counter.

3 Application: a Clinical X-ray Generator

In this section we introduce an industrial control component that we use for reporting our experiences (Section 3.1). We also informally describe a DSL for specifying the behavior of the component (Section 3.2).

3.1 Platform

Philips Healthcare produces interventional X-ray systems (see Fig. 2a) which are used to perform minimally-invasive medical procedures. During a procedure, the surgeon uses the images on the monitors as guidance. Images are constructed for two projections. The X-ray system of Fig. 2a consists of two planes: frontal

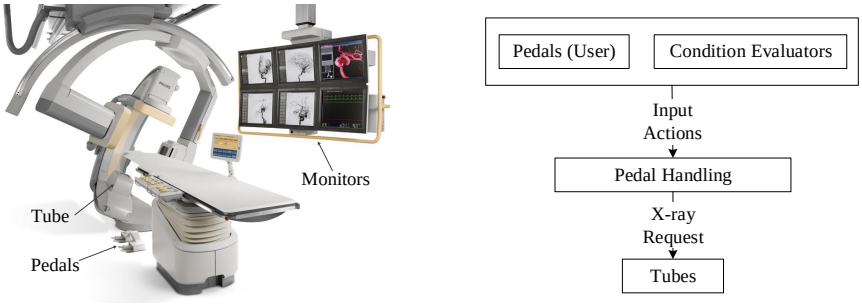

(a) Interventional X-ray

(b) Interfaces of *Pedal Handling*

Figure 2: Industrial Study Case

(top-down) and lateral (left-right). These planes can be used separately or together (biplane).

The surgeon can send X-ray requests using the pedals. The interventional system of Fig. 2a includes a component called *Pedal Handling* that makes decisions about the amount of X-ray that should be generated in the tube of each plane (Fig. 2b). From each plane the following types of X-ray can be generated:

- Fluoroscopy: low dose X-ray, for obtaining real-time images;
- SingleShot: high dose X-ray, for capturing a single image;
- Series: high dose X-ray, for recording a series of images.

Pedal Handling also takes into account conditions that should interrupt the X-ray, or that should prevent the X-ray from starting. *Condition Evaluators* continuously evaluate these conditions and notify *Pedal Handling* when changes occur.

3.2 DSL for Pedal Handling

To describe the behavior of *Pedal Handling*, domain experts mainly focus on the external interfaces of the component. Starting from the initial state, they think about input actions received from *Pedals* and *Condition Evaluators*. Based on the received input, the component might change its current state; it also makes a decision about the output X-ray and sends a request to the tubes. This process continues by receiving the next input. Thus, from the domain expert's point of view the behavior of *Pedal Handling* can be described with alternating sequences of input and output actions.

To specify the behavior of *Pedal Handling*, we use a DSL that fits this way of reasoning. We have used Xtext [4] to define the grammar (concrete syntax) of this DSL. Based on the grammar a parser, an Eclipse-based editor, and a set

of Java classes describing the concepts in the abstract syntax of the language are automatically generated.

Fig. 3 depicts a simplified specification in the DSL[1]. A DSL model starts with declaring the input actions that can be received by *Pedal Handling* (`InActions`). Afterwards, it declares variables that keep track of the current state of the component (`Boolean variables` and `Plane variables`) and their initial values (`Init`). Two not-explicitly-declared variables, i.e., `OutputType` and `OutputPlane`, determine the output of *Pedal Handling*.

The internal logic of the component is described in terms of `Rule`s. Each rule refers to an input action and consists of a guard and a do clause. The guard describes when the input action is enabled and the do clause determines how the action influences the state of the component and the output. A DSL model specifies exactly one rule for each input action. Multiple rules for an action are not supported.

Rules of the DSL use very simple constructs for describing behavior. However, the precise meaning of some constructs is not immediately obvious. For instance, it is not obvious whether evaluating the do clause of an action can be interrupted by receiving a new input action. Since a do clause may contain multiple assignments to a variable, it is also relevant to determine when a variable assignment takes effect.

In Section 4 we identify the main concepts in the DSL (abstract syntax) and give a formalization which clearly specifies the semantics of the language. For instance, our semantics enforces that do clauses should be interpreted in an atomic way and each assignment to a variable immediately changes its value such that the previous value is overwritten.

4 Abstract Syntax and Semantics of the DSL

In Section 3.2 we informally introduced a DSL and motivated the need for a formal semantics that describes the meaning of the language constructs. In this section, we identify the main concepts of the language and give a precise meaning to them by means of a transformation to a language, namely mCRL2, that has a formal semantics. The semantics of the mCRL2 language is defined in [11] using Structural Operational Semantics (SOS) [18].

We use the model from Fig. 3 as a running example to describe our transformation to mCRL2. Our choice of mCRL2 is motivated by the expressiveness of the language, the availability of a toolset [2] that supports analysis of behavior, and our previous experience [14] with the language and toolset.

We first discuss the required data specifications in Section 4.1. Then we use process expressions to describe the behavior specified by DSL models in Section 4.2. In Section 4.1 and 4.2, we have dedicated separate paragraphs to discuss the concepts from the abstract syntax of the DSL. Bold text at the start of each paragraph indicates the name of the concepts.

[1]For confidentiality reasons, details have been changed in this example.

```
InActions: FRFluoOn FRFluoOff
           StartCond ResetStartCond

Boolean variables: FRFluoReq FRFluoOK
Plane variables: FluoPlane

Init: FRFluoReq := false ;
      FRFluoOK := true ;
      FluoPlane := NONE ;

Rule: FRFluoOn guard: FRFluoReq == false
    do: FRFluoReq := true ;
        FluoPlane := FR ;
        if FRFluoOK
            then OutputType := Fluo ; OutputPlane := FR ;
        fi

Rule: FRFluoOff guard: FRFluoReq == true
    do: FRFluoReq := false ;
        if FluoPlane == FR
            then FluoPlane := NONE ;
            if OutputType == Fluo
                then OutputType := Standby ; OutputPlane := NONE ;
            fi
        fi

Rule: StartCond guard: FRFluoOK == true
    do: FRFluoOK := false ;

Rule: ResetStartCond guard: FRFluoOK == false
    do: FRFluoOK := true ;
```

Figure 3: Snapshot of a DSL Model

Finally, in Section 4.3 we discuss a few correctness properties that we verify against DSL models.

4.1 Data Specification

Plane, X-Ray Type As mentioned in Section 3.1, X-ray can be generated from three planes (frontal, lateral, and biplane). We define the data types *Plane* and *XRay* to describe the planes and the type of X-ray generated from them.

sort $\quad Plane = \textbf{struct}\ None \mid FR \mid LT \mid BI;$
$\qquad XRay = \textbf{struct}\ Standby \mid Fluo \mid SingleShot \mid Series;$

The combination of *None* and *Standby* describes a situation in which no X-ray is generated from the planes.

State A DSL model declares a set of boolean and plane variables. The valuation of these variables and the two special variables `OutputType` and `OutputPlane` determine the state of the transition system described by the DSL model.

To describe the notion of state in mCRL2, we create two structured data types B (for boolean variables) and P (for plane variables). The structure of

these sorts corresponds to the variable declarations of the DSL model. For the model of Fig. 3, B and P are defined as follows:

$$\begin{array}{ll}\textbf{sort} & B \;=\; \textbf{struct}\;\; \mathit{FRFluoReq} \mid \mathit{FRFluoOK};\\ & P \;=\; \textbf{struct}\;\; \mathit{FluoPlane};\end{array}$$

To specify the valuations of boolean and plane variables, we declare the following data types:

$$\begin{array}{ll}\textbf{sort} & \mathit{BVals} \;=\; B \to \mathbb{B};\\ & \mathit{PVals} \;=\; P \to \mathit{Plane};\end{array}$$

Finally, the notion of state can be formalized as follows:

$$\textbf{sort}\quad \mathit{PSt} \;=\; \textbf{struct}\;\; \mathit{St}(\mathit{bs}{:}\mathit{BVals}, \mathit{ps}{:}\mathit{PVals}, \mathit{outType}{:}\mathit{XRay}, \mathit{outPlane}{:}\mathit{Plane});$$

The projection functions *outType* and *outPlane* extract the values of `OutputType` and `OutputPlane` from states.

In a DSL model, the initial values of the boolean and plane variables are specified by `Init`. We declare $\mathit{bs}_0 \in \mathit{BVals}$ and $\mathit{ps}_0 \in \mathit{PVals}$ to specify the initial values in mCRL2. We also declare $s_0 \in \mathit{PSt}$ to describe the initial state.

$$\begin{array}{ll}\textbf{map} & \mathit{bs}_0 : \mathit{BVals};\\ & \mathit{ps}_0 : \mathit{PVals};\\ & s_0 : \mathit{PSt};\\ \textbf{eqn} & \mathit{bs}_0(\mathit{FRFluoReq}) = \mathit{false};\\ & \mathit{bs}_0(\mathit{FRFluoOK}) = \mathit{true};\\ & \mathit{ps}_0(\mathit{FluoPlane}) = \mathit{None};\\ & s_0 = \mathit{St}(\mathit{bs}_0, \mathit{ps}_0, \mathit{Standby}, \mathit{None});\end{array}$$

The initial values of `OutputType` and `OutputPlane` cannot be specified by `Init` in the DSL. It is assumed that initially no X-ray is generated from the planes. Hence, we specify s_0 such that:

$$\mathit{outType}(s_0) = \mathit{Standby} \qquad \mathit{outPlane}(s_0) = \mathit{None}$$

Guard, Do Clause, Rule A DSL model specifies one rule for each input action. The rule of an input action consists of a guard and a do clause. From Fig. 3 one can see that a guard is a function from states to booleans. A do clause consists of a sequence of assignments/conditionals that given the current state can change the values of the variables and produce a new state. We describe guards, do clauses, and rules as follows:

$$\begin{array}{ll}\textbf{sort} & \mathit{Guard} \;=\; \mathit{PSt} \to \mathbb{B};\\ & \mathit{DCl} \;=\; \mathit{List}(\mathit{PSt} \to \mathit{PSt});\\ & \mathit{Rule} \;=\; \textbf{struct}\;\; R(\mathit{gr}{:}\mathit{Guard}, \mathit{dc}{:}\mathit{DCl});\end{array}$$

The projection functions *gr* and *dc* extract the guard and do clause element of a rule, respectively.

For the DSL model of Fig. 3 with 4 rules, we declare the guards, do clauses, and rules g_i, d_i, r_i for $1 \leq i \leq 4$:

map $g_1, g_2, g_3, g_4 : Guard$;
$d_1, d_2, d_3, d_4 : DCl$;
$r_1, r_2, r_3, r_4 : Rule$;

The calculations of each guard can be described in terms of an equation. For instance, the guard of the first rule of Fig. 3 can be defined as follows:

var $b : BVals$;
$p : PVals$;
$xr : XRay$;
$pl : Plane$;
eqn $g_1(St(b, p, xr, pl)) = (b(FRFluoReq) \approx false)$;

To describe do clauses, we specify assignments and conditionals in terms of equations. To explain this, we consider the do clause of the first rule from Fig. 3. This do clause contains four assignments and one conditional.

We describe each assignment as a function that updates one of the components of a state $St(b, p, xr, pl)$. We denote the assignments of the first rule by a_1, a_2, a_3, a_4 based on their order of appearance:

map $a_1, a_2, a_3, a_4 : PSt \to PSt$;
eqn $a_1(St(b, p, xr, pl)) = St(b[FRFluoReq \to true], p, xr, pl)$;
$a_2(St(b, p, xr, pl)) = St(b, p[FluoPlane \to FR], xr, pl)$;
$a_3(St(b, p, xr, pl)) = St(b, p, Fluo, pl)$;
$a_4(St(b, p, xr, pl)) = St(b, p, xr, FR)$;

For example, the first assignment updates the values of boolean variables by setting *FRFluoReq* to *true*.

A conditional statement of a do clause is specified by a term of the shape $if(c, t, u)$ in equations. The conditional in the first rule of Fig. 3 can be described as follows:

map $cond : PSt \to PSt$;
eqn $cond(St(b, p, xr, pl)) = if(b(FRFluoOK), Eval(condthen, St(b, p, xr, pl)), St(b, p, xr, pl))$;

In this equation, *Eval* is a function that evaluates a sequence of assignments or conditionals and *condthen* is the "then" part of the conditional (see below).

A do clause is described as a sequence of assignments/conditionals. For example, we describe the do clause of the first rule (d_1) as a sequence of the assignments a_1, a_2 and the conditional *cond*. The "then" part of a conditional is also specified as a sequence of its components. The "then" part of the conditional in the first rule (*condthen*) is a sequence of a_3, a_4:

map $condthen : List(PSt \to PSt)$;
eqn $d_1 = [a_1, a_2, cond]$;
$condthen = [a_3, a_4]$;

Having sequences of assignments and conditionals, it is also essential to define a function that applies a sequence of statements to a state and returns the resulting state. The following mCRL2 description defines *Eval* for this purpose:

$$
\begin{array}{ll}
\textbf{map} & \mathit{Eval} : \mathit{List}(\mathit{PSt} \to \mathit{PSt}) \times \mathit{PSt} \to \mathit{PSt}; \\
\textbf{var} & s : \mathit{PSt}; \\
 & f : \mathit{PSt} \to \mathit{PSt}; \\
 & \ell : \mathit{List}(\mathit{PSt} \to \mathit{PSt}); \\
\textbf{eqn} & \mathit{Eval}([], s) = s; \\
 & \mathit{Eval}(f \triangleright \ell, s) = \mathit{Eval}(\ell, f(s));
\end{array}
$$

The function *Eval* is defined by specifying its effect on terms of the shape [] and $f \triangleright \ell$ (the constructors of *List*).

Finally, each rule consists of a guard and a do clause. For example, the first rule from Fig. 3 (declared as r_1 above) can be defined as follows:

$$
\begin{array}{ll}
\textbf{eqn} & r_1 = R(g_1, d_1);
\end{array}
$$

4.2 Process Specification

Action The *Pedal Handling* component performs two types of actions: input and output. A DSL model explicitly declares a set of input actions (from *Pedals* and *Condition Evaluators*) by `InActions`. *Pedal Handling* also performs actions $\mathit{output}(xr, p)$ to send requests for $xr \in \mathit{XRay}$ to $p \in \mathit{Plane}$. This action is not explicitly declared in the DSL. Corresponding to the input actions and the output action we declare actions in mCRL2. For Fig. 3 we declare:

$$
\begin{array}{ll}
\textbf{act} & \mathit{FRFluoOn}, \mathit{FRFluoOff}, \mathit{StartCond}, \mathit{ResetStartCond}; \\
 & \mathit{output} : \mathit{XRay} \times \mathit{Plane};
\end{array}
$$

In Section 3.2, we mentioned that domain experts describe the behavior of *Pedal Handling* with alternating sequences of input and output actions. The semantics of the DSL is aligned with this intuition. We specify the semantics of the DSL model of Fig. 3 in terms of the following processes:

$$
\begin{array}{ll}
\textbf{proc} & P_{In}(s{:}\mathit{PSt}) = \mathit{gr}(r_1)(s) \to (\mathit{FRFluoOn}.P_{Out}(\mathit{Eval}(\mathit{dc}(r_1), s))) \\
 & \qquad + \mathit{gr}(r_2)(s) \to (\mathit{FRFluoOff}.P_{Out}(\mathit{Eval}(\mathit{dc}(r_2), s))) \\
 & \qquad + \mathit{gr}(r_3)(s) \to (\mathit{StartCond}.P_{Out}(\mathit{Eval}(\mathit{dc}(r_3), s))) \\
 & \qquad + \mathit{gr}(r_4)(s) \to (\mathit{ResetStartCond}.P_{Out}(\mathit{Eval}(\mathit{dc}(r_4), s))); \\
 & P_{Out}(s{:}\mathit{PSt}) = \mathit{output}(\mathit{outType}(s), \mathit{outPlane}(s)).P_{In}(s); \\
\textbf{init} & P_{In}(s_0);
\end{array}
$$

The process P_{In} describes the behavior of *Pedal Handling* when the component is ready to receive an input. It carries a data parameter that indicates the current state. The process P_{In} uses a combination of choices and conditional operators for case distinction. The guards (extracted by $\mathit{gr}(r_i)$) are used as conditions in the conditional operators to determine enabled actions. Performing an input action updates the state based on the corresponding do clause (extracted by $\mathit{dc}(r_i)$).

The process P_{Out} describes the behavior of *Pedal Handling* when the component is ready to produce an output. In this situation, *output* is performed and it carries the X-ray type and plane extracted from the current state. Performing the output action does not influence the state.

The processes P_{In} and P_{Out} enforce alternating execution of the input and output actions. Do clauses are evaluated by *Eval*. Thus, new actions cannot be performed before a do clause is completely evaluated. Moreover, assignments have an immediate effect.

4.3 Validating DSL Models

For a safety critical component, it is desired to use DSL models as a single source to automatically obtain models that enable analysis using various formal techniques.

To facilitate formal analysis and verification, we have automated the transformation from the DSL to mCRL2. We have used the mCRL2 toolset to generate the state spaces of DSL models and to verify properties expressed in a variant of the modal μ-calculus [15]. An actual DSL model declares 25 input actions and their effects in terms of rules. The corresponding state space consists of approximately 45000 states and 350000 transitions. We have verified some safety properties against this model, e.g., "deadlock-freedom", "no X-ray is generated from the planes when there is no request from the user". A formalization of some safety properties for the DSL model of Fig. 3 are included in Appendix A.

The mCRL2 formalization of the semantics can also be used as a reference to implement transformations to formalisms that support other kinds of analysis. In Section 5 and 6 we discuss automated transformations from the DSL to formalisms that enable simulation and model-based testing, respectively. Fig. 4 depicts an overview of the transformations from the DSL that we discuss throughout this paper. We have used Xtend [3] to develop these transformations.

Using the DSL semantics as guidance for implementing such transformations does not give a robust connection between the generated models. However, it helps to avoid obvious inconsistencies between them. In Section 7, we formally validate the consistency between the generated models.

A proposed implementation for *Pedal Handling* is already available. Thus, at the moment we do not generate implementation code from the DSL.

5 Obtaining Simulation Models from DSL Models

In this section we use the semantics of Section 4 as a reference to construct a transformation from the DSL to models that enable interactive simulation of behavior. We aim to use models in the parallel object-oriented specification language [23] (POOSL) for simulation. Tool support [19] and our previous experience with the language motivated our choice.

POOSL is a light-weight modeling language and its formal semantics is expressed in terms of timed probabilistic labeled transition systems [26]. The tools available for POOSL allow us to simulate models in the POOSL language

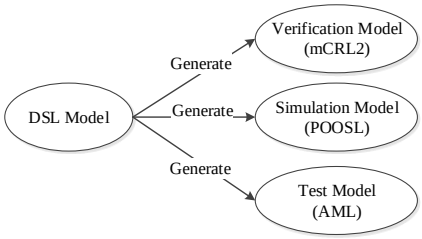

Figure 4: Generated Artifacts for *Pedal Handling*

and discuss the modeled behavior with domain experts. Thus, it is desired to have an automated transformation that generates simulation models from DSL specifications.

In this section we first provide a brief overview of the POOSL language (Section 5.1). Afterwards, we discuss our transformation from the DSL to simulation models in POOSL (Section 5.2).

5.1 An Overview of POOSL

In this section we discuss a few constructs of the POOSL language; our focus is on the syntactic constructs that we need for realizing the semantics of Section 4 in POOSL. The interested reader can refer to [26] and [19] for more details about the language and some instructive examples. We use the modulo 3 counter from Section 2 to illustrate the relevant constructs of the language.

In POOSL behavior is described in terms of process classes. Process classes are templates that need to be instantiated to obtain concrete instances that are executable. Each process class declares a set of communication ports that are available to its instances for sending/receiving messages. Process classes also define the signature of messages that can travel across the declared ports. For instance, Fig. 5 depicts a process class called `Counter` with two ports, namely, `cin` and `cout`. The message declarations indicate that messages `count` and `reset` can be communicated through the ports `cout` and `cin`, respectively. The message `count` carries an integer parameter to indicate the current number.

In POOSL there is a notion of initiative for messages, whereas in mCRL2 we do not explicitly define the direction of interactions for actions. The symbols `!` (message send) and `?` (message receive) denote the direction of communication for each message. For our simple example, we choose to declare `count` as a message sent by `Counter` and `reset` as a message that can be received by `Counter`.

Process classes also consist of variables. Variables of a process are only accessible by the process itself and keep track of data items that are relevant for describing the behavior of the process. Declaring a variable involves specifying a name and its data type. The process `Counter` in Fig. 5 declares variable `n` of

```
/* Counter Process */
process class Counter()
ports
    cin,cout
messages
    cout ! count(Integer),
    cin ? reset
variables
     n : Integer
init
    initialize()()

methods
initialize()()
    n := 0 ;
    countMethod()()

countMethod()()
    sel
        [(n < 3)] cout ! count(n);
            n := n+1;
            countMethod()()
        or
        [(n == 3)] cin ? reset;
            n := 0;
            countMethod()()
    les
```

Figure 5: A Modulo 3 Counter in POOSL

type integer; this variable is used to store the current number.

The behavior of a process class is specified in terms of its methods. Each process has an initial method specified by `init`. Generally speaking, the behavior defined by a method is specified by the following syntax:

$$\texttt{m}(\texttt{p}_1\texttt{:d}_{\texttt{p}_1},\ldots,\texttt{p}_k\texttt{:d}_{\texttt{p}_k})(\texttt{r}_1\texttt{:d}_{\texttt{r}_1},\ldots,\texttt{r}_m\texttt{:d}_{\texttt{r}_m})\ |\texttt{l}_1\texttt{:d}_{\texttt{l}_1},\ldots,\texttt{l}_n\texttt{:d}_{\texttt{l}_n}|$$
$$\texttt{S}$$

where $\texttt{p}_i$ (for $1 \leq i \leq k$), $\texttt{r}_i$ (for $1 \leq i \leq m$), and $\texttt{l}_i$ (for $1 \leq i \leq n$) are the input, return, and local parameters of the method, respectively. For each parameter, the corresponding data type is also specified.

The body of a method (denoted by `S` in the example above) is composed of different types of statements. In what follows, we discuss a few types of statements that can appear inside a method. Suppose `x` is a process variable, `E` is an expression, `p` is a port, $\texttt{S},\texttt{S}_1,\texttt{S}_2$ are statements, and `C` is a boolean expression. The statement `x:=E` evaluates the expression `E` and assigns its value to `x`. Here we assume that `E` is a well-defined expression in terms of process variables, constants from standard data types (such as natural numbers and booleans), and common operations on these data types (such as addition for integers); see [26] for a detailed description of POOSL's data layer and the syntax of expressions.

The statement `p!m` denotes a message `m` being sent to the port `p`. In POOSL communication is synchronous, i.e., a send statement `p!m` is blocked until a

```
/* User Process */
process class User()
ports
    uin,uout
messages
    uin ? count(Integer),
    uout ! reset
variables
init
    observeMethod()()

methods

observeMethod()()| X : Integer |
    sel
        uin ? count(X);
            observeMethod()()
        or
        uout ! reset;
            observeMethod()()
    les
```

Figure 6: A Process Specification for Communicating with Instances of `Counter`

matching receive statement `q?m` can be executed such that the ports `p` and `q` are connected via a channel; see description below about connecting ports with channels.

The basic statements mentioned above can be combined in different ways to obtain more complex behaviors. The non-deterministic choice between S_1 and S_2 is denoted by `sel` S_1 `or` S_2 `les`. The notation $S_1;S_2$ denotes the sequential composition of S_1 and S_2.

The guarded command `[C]S` evaluates the condition `C` and only if `C` evaluates to *true* the statement `S` is allowed to execute. Otherwise, deadlock will occur. The statement `if C then` S_1 `else` S_2 `fi` results in executing S_1 if `C` evaluates to *true* and otherwise it executes the statement S_2.

The process specification in Fig. 5 declares two methods, namely, `initialize` and `countMethod`; instances of `Counter` initially behave as `initialize`. The method `initialize` sets `n` to zero and calls `countMethod`. A combination of non-deterministic choices and guarded commands are used for case distinction in `countMethod`. If `n<3` holds, the process sends a `count` message that carries the current number. Afterwards, it increments the current number by one and

```
/* System Specification */
system
instances
    c : Counter()
    u : User()
channels
    {c.cout , u.uin}
    {c.cin , u.uout}
```

Figure 7: A Communication Channel between Instances of `Counter` and `User`

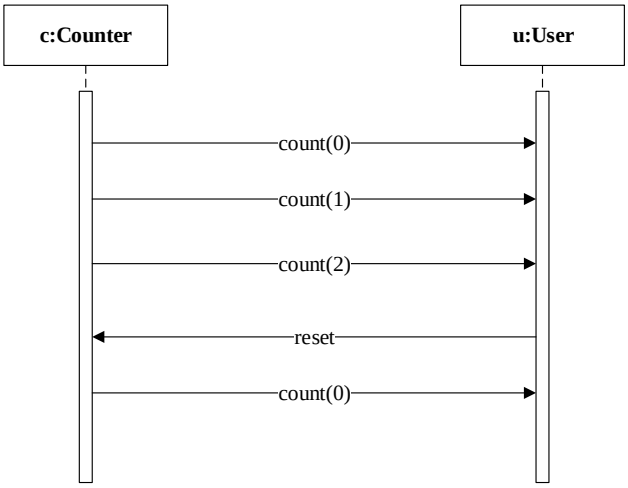

Figure 8: Simulating the Interactions between a `Counter` and `User`

calls `countMethod` to continue the counting. If `n==3` holds, the process receives `reset` and resets the current number to zero. The method `countMethod` is called to restart the counting.

As indicated earlier, sending and receiving messages in POOSL should be performed in a synchronous way. For instance, in our example, `cout!count(n)` is only executable, if another process is ready to execute a complementary message, e.g., `uin?count(X)`, such that `uin` is connected to `cout` via a channel. By synchronizing these two messages the value of `n` is bound to `X`. Similarly, executing `cin?reset` require a complementary send message from another process. To this end, we specify a process class called `User` (see Fig. 6) to exchange messages with `Counter`.

The behavior of `User` is described in terms a method called `observeMethod`. The non-deterministic choice in `observeMethod` allows instances of the process to receive `count` and send `reset` when the complementary messages are available.

Finally, to obtain an executable model for the modulo 3 counter, we need concrete instances of `Counter` and `User` which are connected via suitable channels. The system specification from Fig. 7 instantiates the defined process classes and creates two communication channels between the instances by connecting their ports. Fig. 8 illustrates a sequence diagram obtained by simulating the resulting model; the diagram consists of two lifelines (for the instances `c` and `u` of the process classes) and depicts the first five steps of their communication.

5.2 A Transformation from the DSL to POOSL

In this section we describe our transformation scheme for generating POOSL models from DSL models. Similar to the example model from Section 5.1, simulation models generated from the DSL consist of two process classes:

- `DecisionClass`: receives the inputs declared in the DSL model and makes decisions about the output X-ray based on the rules specified in the DSL model;
- `User`: sends the inputs declared in the DSL model to `DecisionClass` and observes the produced outputs.

In Section 5.1, we specified a simple process class for observing the behavior performed by instances of `Counter`. A similar structure can be used to specify the process class `User` to observe the behavior described by a given DSL model. Thus, in this section we focus on the way we obtain `DecisionClass` from a given DSL model. We use the DSL model from Fig. 3 to illustrate the way the domain concepts identified in Section 4 are described in POOSL.

Action In our transformation, port and message declarations of `DecisionClass` are based on the distinction between the inputs and outputs of *Pedal Handling*. We declare two ports, namely, `in` (for receiving inputs) and `out` (for sending outputs); see the port and message declarations of Fig. 9 for the DSL model from Fig. 3.

Messages that are communicated through the port `in` correspond to the input actions declared by `InActions` in a DSL model. For the DSL model from Fig. 3 we declare `FRFluoOn,FRFluoOff,StartCond,ResetStartCond` as messages that travel across the port `in`; see Fig. 9.

As indicated earlier, the output actions of *Pedal Handling* are not explicitly declared in the DSL. The process `DecisionClass` sends `Output` on the port `out` to model the outputs performed by *Pedal Handling*. In comparison with the mCRL2 formalization of Section 4, we do not use custom data types in POOSL to describe the types of X-ray and the planes; see the declaration of *XRay* and *Plane* in Section 4. Using custom data types in POOSL requires defining new data classes and manipulating their instances. To avoid such complexities in our transformations, we declare `Output` as an action with two string parameters (see Fig. 9); performing `Output("Fluo","FR")` indicates a request for X-ray type *Fluo* to plane *FR*. This restricted way of using string parameters is consistent with the use of *XRay* and *Plane* as the types of parameters carried by output actions in mCRL2.

Plane, X-Ray Type In Section 4 we defined custom data types to describe the planes (i.e, *None*, *FR*, *LT*, *BI*) and the X-ray types (i.e., *Standby*, *Fluo*, *SingleShot*, *Series*). As mentioned above, for communicating outputs, we use strings to represent these elements. Since in object-oriented languages performing basic operations on strings, e.g., "equality checking", often leads to confusing

```
ports
    in, out
messages
/* input messages */
    in ? FRFluoOn,
    in ? FRFluoOff,
    in ? StartCond,
    in ? ResetStartCond,
/* output message */
    out ! Output(String,String)
```

Figure 9: Declaring Input and Output Messages in POOSL for the Example DSL Model

interpretations such as "checking equality of values" and "checking equality of referenced objects", we prefer to have a simple way of representing the planes and X-ray types in the internal computations of DSL rules. We represent the elements of these two domain concepts by the sets $\{1, 2, 3, 4\}$ and $\{5, 6, 7, 8\}$, respectively. Fig. 10 depicts the mapping we apply for representing the planes and X-ray types.

To represent this mapping in POOSL, we declare an integer variable for each element depicted in the tables of Fig. 10; see the variable declarations of Fig. 11. These variables are initialized based on the defined mapping; the initial values are specified by the method `initialize` in Fig. 11.

To ensure that the mapping is correctly used, the values of these variables are not changed by the methods of the process class after initialization.

State As discussed in Section 4, the notion of state in the DSL is defined based on the set of variables declared in a DSL model (boolean and plane variables) and the two not-explicitly declared output variables (`OutputType` and `OutputPlane`). The declaration of boolean variables in POOSL corresponds to the boolean variables of the DSL model; see the declarations of boolean variables in Fig. 11 for the DSL model from Fig. 3. Plane variables and the non-explicitly declared output variables (`OutputType` and `OutputPlane`) are denoted by integer variables in POOSL; see the declarations of integer variables in Fig. 11 for the DSL model from Fig. 3.

The initial values of the boolean and plane variables are specified in the `initialize` method and correspond to the initial values described in the DSL model. The output variables are initialized based on the assumption that in the

Plane	**Integer Value**
None	1
FR	2
LT	3
BI	4

X-ray Type	**Integer Value**
Standby	5
Fluo	6
SingleShot	7
Series	8

Figure 10: Representing Planes and X-ray Types in POOSL

```
variables
/* mapping variables */
    None : Integer ,
    FR : Integer ,
    LT : Integer ,
    BI : Integer ,
    Standby : Integer ,
    Fluo : Integer ,
    SingleShot : Integer ,
    Series : Integer ,
/* boolean variables */
    FRFluoReq : Boolean ,
    FRFluoOK : Boolean ,
/* plane variables */
    FluoPlane : Integer,
/* output variables */
    OutputType : Integer,
    OutputPlane : Integer
init
    initialize()()
methods

initialize()()
/* mapping variables */
    None := 1;
    FR := 2;
    LT := 3;
    BI := 4;
    Standby := 5;
    Fluo := 6;
    SingleShot := 7;
    Series := 8;
/* boolean variables */
    FRFluoReq := false;
    FRFluoOK := true;
/* plane variables */
    FluoPlane := None;
/* output variables */
    OutputType := Standby;
    OutputPlane := None;

    inputMethod()()
```

Figure 11: Declaring State Variables in POOSL for the Example DSL Model

initial state no X-ray is generated from the planes.

Guard, Do Clause, Rule A DSL rule refers to an input action and consists of a guard (specifying when the action is active) and a do clause (specifying the effects of performing the input action). We describe each rule by a guarded command `[C]S` where `C` is the guard of the DSL rule and `S` describes the do clause. Fig. 12 depicts a specification of the first rule from Fig. 3.

The do clause of a DSL rule is a sequence of assignments/conditionals. As discussed in Section 5.1, `x:=E` and `if C then` S_1 `else` S_2 `fi` are among the statement types that can appear in the body of a POOSL method. Thus, we use these constructs to specify do clauses in POOSL. The do clause of the first rule from Fig. 3 is translated to four assignments and one conditional in POOSL;

```
[(FRFluoReq == false)] in ? FRFluoOn ;
        FRFluoReq := true;
        FluoPlane := FR;
        if FRFluoOK
            then
                OutputType := Fluo;
                OutputPlane := FR
        fi;
```

Figure 12: Describing a Rule from the Example DSL Model in POOSL

see the statements guarded by `[(FRFluoReq==false)]` in Fig. 12.

We repeat this transformation scheme for all the rules of a DSL model and specify a method called `inputMethod` that describes the way input actions are received by *Pedal Handling*. Fig. 13 depicts this transformation for the DSL model from Fig. 3. The body of `inputMethod` consists of a non-deterministic choice between the guarded commands that describe the DSL rules. Note that in Fig. 13 we have folded the descriptions of the last three rules.

Observe that in Fig. 13 we have extended the description of Fig. 12 for the first rule of the DSL model. A sequential composition of the do clause and `outputMethod` is used to show that after computing the effects of the action, *Pedal Handling* is ready to produce an output. Each folded specification block also consists of a sequential composition of the corresponding do clause and `outputMethod`.

The method `outputMethod` describes the behavior of *Pedal Handling* when the component is ready to perform an output action. A combination of non-deterministic choices and guarded commands are applied to transform the integer values of `OutputType` and `OutputPlane` into string parameters for the action `Output` (see the mapping from Fig. 10). After performing an output action the method `inputMethod` is called and the model can receive the next input.

Similar to the process P_{In} and P_{Out} in the mCRL2 formalization, we use the methods `inputMethod` and `outputMethod` to enforce alternative execution of the input and output actions.

6 Validating the Industrial Implementation by Model-based Testing

In industry, the implementation is considered to be the most valuable artifact that is produced for a component. When behavioral models of a component are also available, it is relevant to check whether the implementation complies to the modeled behavior. Validating the compliance of implementations to DSL models adds a level of redundancy that can reveal discrepancies between the developed DSL models and implementations.

We use model-based testing to validate the correctness of an implementation with respect to a DSL model. Model-based testing uses a model that describes the behavior of a system under test and enables both automatic generation and

```
inputMethod()()
    sel
        [(FRFluoReq == false)] in ? FRFluoOn ;
                FRFluoReq := true;
                FluoPlane := FR;
                if FRFluoOK
                    then
                        OutputType := Fluo;
                        OutputPlane := FR
                fi;
                outputMethod()()
    or
        [(FRFluoReq == true)] in ? FRFluoOff ;
    or
        [(FRFluoOK == true)] in ? StartCond ;
    or
        [(FRFluoOK == false)] in ? ResetStartCond ;
    les

outputMethod()()
    sel
        [(OutputType == Standby) & (OutputPlane == None)] out ! Output("Standby","None")
    or
        [(OutputType == Fluo) & (OutputPlane == FR)] out ! Output("Fluo","FR")
    or
        [(OutputType == Fluo) & (OutputPlane == LT)] out ! Output("Fluo","LT")
    or
        [(OutputType == Fluo) & (OutputPlane == BI)] out ! Output("Fluo","BI")
    or
        [(OutputType == SingleShot) & (OutputPlane == FR)] out ! Output("SingleShot","FR")
    or
        [(OutputType == SingleShot) & (OutputPlane == LT)] out ! Output("SingleShot","LT")
    or
        [(OutputType == SingleShot) & (OutputPlane == BI)] out ! Output("SingleShot","BI")
    or
        [(OutputType == Series) & (OutputPlane == FR)] out ! Output("Series","FR")
    or
        [(OutputType == Series) & (OutputPlane == LT)] out ! Output("Series","LT")
    or
        [(OutputType == Series) & (OutputPlane == BI)] out ! Output("Series","BI")
    les;
    inputMethod()()
```

Figure 13: POOSL Methods for Receiving Inputs and Producing Outputs

execution of test cases on the implementation.

The semantics of the pedal handling DSL is described in terms of labeled transition systems and hence we use the theory of input-output conformance (*ioco*) testing for labeled transition systems [24] to automatically derive test cases from the behavior described in the DSL and execute them on the implementation.

In the *ioco* theory, correctness of implementations with respect to specifications is expressed in terms of the binary relation of *ioco*. The *ioco* theory provides an algorithm that derives a set of test cases from a given specification such that executing this set of test cases on an implementation determines whether the specification and implementation are related by the conformance relation.

The model-based testing tool of Axini [1] is based on the *ioco* theory. In this tool, specifications are described using the Axini Modeling Language (AML). In

order to facilitate model-based testing on implementations of the *Pedal Handling* component, we aim to develop an automated transformation from the DSL to AML (Fig. 4).

In this section first we introduce the AML language (Section 6.1) and describe our transformation from the DSL to AML (Section 6.2). Then we discuss about model-based testing of the industrial implementation using the generated models and interpreting the obtained results (Section 6.3). Finally, we discuss some practical remarks (Section 6.4).

6.1 An Overview of AML

AML is a modeling language that allows the modeler to formally specify the behavior of a system that communicates with its environment through sending and receiving actions. Symbolic labeled transition systems is the formalism used by Axini for giving semantics to the language constructs. In what follows, we give a brief overview of the syntax of AML[2]; our focus is on the language constructs that we need throughout this paper for constructing test models from specifications in the *Pedal Handling* DSL. We use the modulo 3 counter from Section 4 as a running example to discuss the constructs of AML.

In AML behavior is described in terms of processes. An AML process starts by declaring the inputs actions (stimuli) that it can accept and the observable output actions (responses) that it can produce. For instance, the following AML statements declare a stimulus (called `reset`) and a response (called `count`) for the modulo 3 counter:

```
stimulus 'reset'
response 'count' , 'curr':integer
```

For each action, the direction of communication should be denoted by indicating the keyword `stimulus` (input) or `response` (output). Actions can carry data parameters. For instance, in the example above `count` carries an integer parameter called `curr` to indicate the current number.

Action declarations in AML processes is followed by declaring state variables. State variables of a process are data items that determine the current state of the process and can influence the course of actions. AML supports various data types for variables, e.g., booleans, integer numbers, strings. A variable declaration in AML consists of a name, a data type, and the initial value of the variable. For the modulo 3 counter, we declare an integer variable called `n` to keep track of the current number. The following AML statement declares this variable and initializes it to zero:

```
var 'n', :integer, 0
```

[2]AML is a modeling language designed by Axini and is only available as part of their commercial model-based testing toolset. For confidentiality reasons, we use a modified version of the language in this paper.

Assigning values to state variables and receiving/sending actions are among the basic operations that can be performed by a process. The following list provides examples of these basic operations for the modulo 3 counter:

- `update 'n=0'`: resetting the state variable `n` to zero;
- `receive 'reset'`: receiving the `reset` action from the environment;
- `send 'count', constraint: 'curr==n'`: sending the `count` action to report the current number to the environment.

The `send` example shows that for actions carrying data parameters we can specify a `constraint`, which is a boolean expression, to describe the desired values that are communicated with the environment.

AML also provides operators that combine the basic operations and allow us to specify more complex behaviors. Suppose S_1,S_2 are AML statements. The non-deterministic choice between S_1 and S_2 is denoted as follows:

```
choice {
     o { S1 }
     o { S2 }
 }
```

A specification that consist of S_1 and S_2 with a line break in between describes the sequential composition of S_1 and S_2.

The values of state variables in a process can also influence the course of actions in the process. Suppose that `S` is an AML statement and `C` is a boolean expression defined in terms of the state variables of a process. The guarded statement `guard 'C' S` only allows the execution of `S` if the condition `C` hold. If `C` evaluates to *false*, the guarded statement `guard 'C' S` will lead to a deadlock.

In Fig. 14, we apply the constructs discussed above to specify the modulo 3 counter in AML. Note that the behavioral description is enclosed in a block labeled by `ready`. This block consists of a non-deterministic choice between two disjoint cases, namely, `n<3` and `n==3`. In the first case, `count` is performed and the current number is incremented by 1, whereas in the second case `reset` is performed and `n` becomes zero. In both cases, the `goto` statement allows us to repeat the behavior defined in the `ready` block.

For a given AML model, the toolset of Axini can generate a symbolic labeled transition system that corresponds to the behavior described in the model. Fig. 15 depicts the behavior of `Counter` in terms a symbolic labeled transition system; to distinguish between the guards and assignments, we put the guards between square brackets.

6.2 A Transformation from the DSL to AML

In this section we consider the domain concepts identified in Section 4 and discuss our transformation scheme for generating AML models from models in the *Pedal Handling* DSL. We use the DSL model from Fig. 3 to describe our transformation.

```
process('Counter') {

    stimulus 'reset'
    response 'count','curr' => :integer

    var 'n', :integer, 0

    label 'ready'
        choice {
         o {
           guard 'n<3'
              send 'count', constraint: 'curr==n'
              update 'n=n+1'
              goto 'ready'
           }
         o {
           guard 'n==3'
              receive 'reset'
              update 'n=0'
              goto 'ready'
           }
    }
}
```

Figure 14: A Process Specification of the Modulo 3 Counter in AML

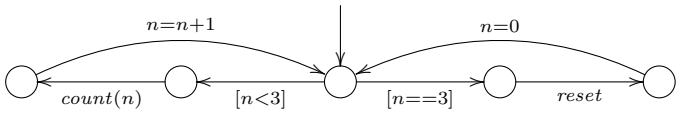

Figure 15: A Symbolic Labeled Transition System Describing the Behavior of `Counter`

Action For a given DSL model, we declare two types of actions in AML:

- stimuli (input actions), correspond to the actions declared by `InActions` in the DSL model;
- response (output action), an action carrying two string parameters that indicate an X-ray request for a specific plane.

For instance, for the DSL model from Fig. 3, this leads to the following action declarations:

```
stimulus 'FRFluoOn','FRFluoOff','StartCond','ResetStartCond'
response 'Output','xraytype'=>:string,'xraychannel' =>:string
```

When sending the action `Output`, we can use the parameter names `xraytype` and `xraychannel` to specify constraints on the data values that are sent through this action.

Plane, X-Ray Type In Section 4 we used custom data types to describe the planes and X-ray types. Although the non-disclosed parts of AML allow us to define custom data types, we have decided to use a limited set of constructs from AML and to represent these two domain concepts with strings.

Strings are a predefined data type in AML. Since AML is not an object-oriented language, accessing strings by reference is not a relevant issue in the language. All operations on strings directly manipulate the values. Thus, we use:

- `"None"`, `"FR"`, `"LT"`, `"BI"` to represent the planes;
- `"Standby"`, `"Fluo"`, `"SingleShot"`, `"Series"` to represent the X-ray types.

State For a given DSL model, the notion of state is defined based on the declared boolean variables, plane variables, and the two non-explicitly declared output variables (`OutputType` and `OutputPlane`). In AML this notion of state is described in terms of booleans (representing the boolean variables of the DSL model) and string variables (representing the plain and output variables of the DSL model).

For the DSL model from Fig. 3 with two boolean variables (i.e., `FRFluoReq` and `FRFluoOK`) and one plane variable (i.e., `FluoPlane`), we use the following declarations in AML (statements preceded by `#` are comments):

```
#boolean variables
var 'FRFluoReq', :boolean, false
var 'FRFluoOK', :boolean, true

#plane variables
var 'FluoPlane', :string, "None"

#output variables
var 'OutputChannel', :string, "None"
var 'OutputType', :string, "Standby"
```

The declared variables are also initialized; the initialization of the boolean and plane variables are based on the initial values specified by `init` in the DSL model (see Fig. 3). The initial values of `OutputChannel` and `OutputType` are based on the domain assumption that initially no X-ray is generated from the planes.

Guard, Do Clause, Rule A DSL model specifies one rule for each input action to indicate when the action is executable (i.e., guard of the rule) and the way performing the action influences the current state (i.e., do clause of the

rule). We transform each rule to a guarded statement `guard 'C' S` in AML such that `C` is the guard of the rule and `S` describes the do clause.

Since do clauses consist of assignments and conditionals, we have to provide transformation schemes for these two constructs. Assignments are transformed to `update` statements in AML. A conditional statement `if C then S fi` from the DSL is transformed to a non-deterministic choice with two options as follows:

```
choice {
   o { guard 'C'  S}
   o { guard '!C'  }
}
```

The first option allows the execution of `S`, if `C` evaluates to *true*. The second option is a guarded statement with an empty body (i.e., skip) and gets executed if `!C` holds.

The following AML statements describe the first rule of the DSL example from Fig. 3:

```
guard '(FRFluoReq == false)'
   receive 'FRFluoOn'
   update 'FRFluoReq=true'            # do clause assignment
   update 'FluoPlane="FR"'            # do clause assignment
   choice {                           # do clause conditional
     o {
       guard 'FRFluoOK'               # then part of conditional
         update 'OutputType="Fluo"'   # do clause assignment
         update 'OutputChannel="FR"'  # do clause assignment
      }
     o {
       guard '!(FRFluoOK)'
      }
   }
```

The do clause of the first rule from Fig. 3 consists of four assignments and one conditional statement. As illustrated above, the assignments are transformed to updates and the conditional statement is described in terms of a choice in AML.

For a given DSL model, we apply this transformation scheme to each rule of the model and use the obtained AML fragments to construct a process specification that describes the behavior specified in the DSL model. Fig. 16 depicts the process specification that we obtain by applying our transformation to the DSL model from Fig. 3; we have folded the specification blocks of the last three rules.

In Fig. 16 the behavioral specification is encloses in a block labeled by `' ready'`. This block consists of a non-deterministic choice between the rules of the DSL model. The guarded statement inside each option describes a rule and only allows the rule to be chosen if it is active (i.e. the guard evaluates to *true*).

```
process('DecisionProcess') {
  stimulus 'FRFluoOn','FRFluoOff','StartCond','ResetStartCond'
  response 'Output','xraytype'=>:string,'xraychannel' =>:string

  #boolean variables
  var 'FRFluoReq', :boolean, false
  var 'FRFluoOK', :boolean, true

  #plane variables
  var 'FluoPlane', :string, "None"

  #output variables
  var 'OutputChannel', :string, "None"
  var 'OutputType', :string, "Standby"

  label 'ready'
  choice {
   o { #Rule 1
     guard '(FRFluoReq == false)'
       receive 'FRFluoOn'
       update 'FRFluoReq=true'         # do clause assignment
       update 'FluoPlane="FR"'         # do clause assignment
       choice {                        # do clause conditional
         o {
          guard 'FRFluoOK'             # then part of conditional
          update 'OutputType="Fluo"'   # do clause assignment
          update 'OutputChannel="FR"'  # do clause assignment
         }
         o {
          guard '!(FRFluoOK)'
         }

         # performing output action
         send 'Output', constraint: 'xraytype==OutputType &&
                                     xraychannel==OutputChannel'
         #return to 'ready'
         goto 'ready'
    }

   o { #Rule 2 }
   o { #Rule 3 }
   o { #Rule 4 }
  }
}
```

Figure 16: An AML Process Specification for the Example DSL Model

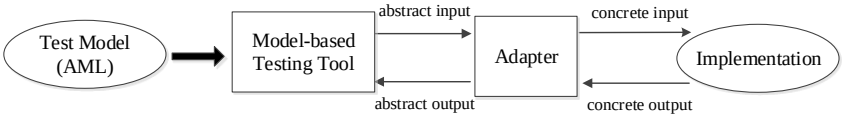

Figure 17: Model-based Testing

The first option in the non-deterministic choice describes the first rule from Fig. 3. Observe that we have extended the specification of the first rule with two sequential compositions. This indicates that after evaluating the do clause an output action is performed and then the behavior of the `'ready'` block is repeated. The folded blocks also have these sequential compositions as their last steps. This enforces the alternating execution of the input and output actions as specified by the processes P_{In} and P_{Out} in Section 4.

6.3 Interpreting the Results of Model-based Testing

Validating the compliance of an implementation to a DSL model by model-based testing requires extra care to interpret the results correctly. A failed test case shows a discrepancy between the behavior described by the test model and the behavior implemented in the implementation. This can have two different reasons:

- a mistake in the test models generated from the DSL;
- a failure in the implementation.

The transformation from the DSL to test models should preserve the semantics of the DSL. A mistake in realizing the semantics may result in failed test cases. In Section 7 we introduce an approach to gain more confidence in the correctness of the generated models in Fig. 4. When there is sufficient confidence in the correctness of test models, a failed test case indicates a failure of the implementation.

It should be noted that if implementations are not generated from the DSL, failed test cases could also indicate a mistake in DSL models; the intended behavior is not correctly described in the DSL but the implementation has correctly realized the behavior.

6.4 Practical Remarks

Fig. 17 depicts the environment we have developed for model-based testing. The tool of Axini uses an on-the-fly algorithm for test case generation and execution. This means that in each step of a test case the tool can do one of the following:

- send an input action allowed by the model to the implementation;

- wait for an output from the implementation and check that the produced output is allowed by the model.

The adapter in Fig. 17 translates the abstract actions from the model to concrete commands for the implementation and maps the concrete outputs produced by the implementation to abstract output actions for the model.

Results In our model-based testing experiments, we encountered failed test cases that indicated a failure in a not-yet-released implementation which is not generated from the DSL. For example, one of the requirements of *Pedal Handling* indicates that high-dose X-ray requests have priority over low-dose X-ray requests. A design mistake in realizing this requirement led to a failed test case. The failed test case was part of an unlikely scenario where three different pedals must be pressed at the same time.

Model-based testing revealed that certain modeling choices taken in our DSL models are implemented differently in the implementation. Unlike stopping *Fluo* (modeled by *FRFluoOff* in the example of Fig. 3), stopping *Series* requires performing two actions in a specific order. The output specified for the first step of stopping *Series* was different from the output produced by the implementation.

7 Checking the Correctness of Generated Models through Redundancy

Transformations from a DSL to analysis models allow the user to apply various formal techniques and reason about DSL models. Analysis models can also be used to assess the correctness of other artifacts available for a component (see Section 6). However, the results obtained from analysis models are only valuable if the corresponding transformations correctly realize the semantics of the language.

Developing a transformation from a DSL to a modeling language requires a deep understanding of the semantics of the DSL and the target language. Moreover, the transformation should not deviate from the semantics of the DSL. Introducing redundancy is a very effective way to reduce the rate of mistakes in such error-prone tasks [27].

We introduce redundancy to validate the correctness of the transformations depicted in Fig. 4 in two ways: equivalence checking (Section 7.1) and model-based testing (Section 7.2).

7.1 Checking the Behavioral Equivalence between Artifacts

The verification and simulation models in Fig. 18 have an underlying labeled transition system. To get confidence in the correctness of simulation models with respect to the formalized semantics (transformation to verification models), we can investigate whether the labeled transition systems described by the

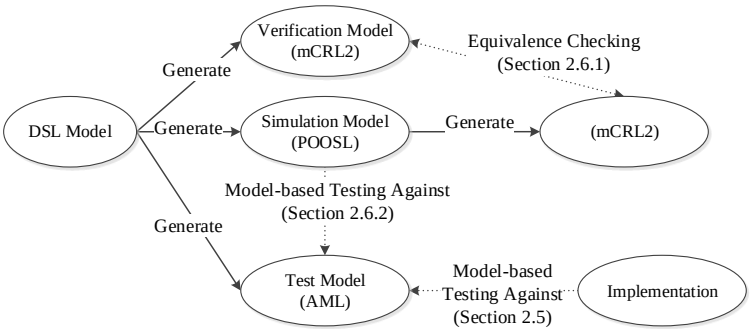

Figure 18: Validation of Artifacts

simulation and verification models are related by an equivalence relation, e.g., strong bisimulation. This may require to develop transformations from analysis models to a formalism that enables state space generation and comparison.

We have used the mCRL2 toolset for state space generation and comparison. To enable state space generation for simulation models, we have used a previously implemented transformation[3] from POOSL to mCRL2; see Section 8 for some observations regarding this transformation.

When comparing behaviors, the internal steps performed by them are not relevant; we focus on the observable behaviors. Hence, we check whether the state spaces are equivalent modulo branching bisimulation. Fig. 18 depicts equivalence checking between mCRL2 and POOSL models and the automated transformation from POOSL to mCRL2 that enables this check.

Results For actual DSL models, the behaviors described in mCRL2 and POOSL were equivalent modulo branching bisimulation. Models generated from three transformations (DSL to mCRL2, DSL to POOSL, and POOSL to mCRL2) are involved in equivalence checking between simulation and verification models. Each transformation is implemented by a different person. This reduces the probability of identical mistakes in the transformations and makes the redundancy introduced by equivalence checking more valuable.

7.2 Model-based Testing of Executable Models

In Section 6 we used model-based testing to validate the correctness of an implementation. Executable analysis models can also be treated as black-boxes that interact via their interfaces; we can supply inputs to an executable model and observe its output. Thus, executable models can be tested against the test

[3]The transformation from POOSL to mCRL2 is based on the formal semantics of POOSL represented in [26] and only supports the syntactic constructs discussed in Section 5. The transformation is implemented by Arjan J. Mooij.

model generated from a DSL specification using model-based testing. Failed test cases reveal mistakes in the transformations.

In our case study, we have applied model-based testing to POOSL models; see Fig. 18. To this end, we have developed the environment of Fig. 17 with a dedicated adapter for POOSL models.

Results We did not encounter any failed test cases in our model-based testing experiments against POOSL models. Similar to Section 7.1, the transformations to simulation models (POOSL) and test models (AML) are implemented by different people. This gives more confidence that the semantics of the DSL is realized correctly in these models.

8 Experiences with Model Transformations

The approach of Fig. 18 deploys model transformations to enable the use of multiple formal techniques and to introduce redundant mechanisms for assessing the correctness of different artifacts with respect to the formalized semantics. This approach relies on two types of transformation: transformations from the DSL to a general-purpose modeling language (from the DSL to mCRL2, POOSL, and AML), and transformations between general-purpose modeling languages (from POOSL to mCRL2). In this section we report on our experiences with these two types of transformations.

General-purpose modeling languages originate from different disciplines and are applicable to a wide range of problems. For instance, POOSL is designed to be applicable for simulation and performance analysis, whereas mCRL2 is focused on formal verification. In our experience, making transformations between two general-purpose formalisms is not a trivial task. There are usually language constructs from the source language that are difficult or even impossible to translate to the target language.

For example, the data layer of POOSL is object-oriented. Each data class describes its variables and methods. Moreover, instances of certain data structures (e.g., strings, lists) are accessed using pointers. On the other hand, mCRL2 is not object-oriented and does not support pointers. Hence, transforming data classes or pointers from POOSL would require complex mechanisms in mCRL2 models which in turn would increase the rate of mistakes in the transformation.

Similarities between the source and target languages may suggest a direct mapping between certain constructs. However, similar constructs in two formalisms may have subtly different semantics. For instance, one would expect the conditional statement `if c then p else q fi` of POOSL to be trace equivalent to $c \rightarrow p \diamond q$ in mCRL2. Based on the formal semantics of POOSL, `if c then p else q fi` first performs an internal step (τ transition) to evaluate the condition. Then it behaves as `p` if `c` holds and otherwise it behaves as `q`. However, in mCRL2 no internal action is performed for evaluating c and hence the two conditional statements are not trace equivalent.

As another example, we consider `[c]p` in POOSL and $c \rightarrow p$ in mCRL2. In both cases, the condition c is evaluated without performing an internal step. For both `[c]p` and $c \rightarrow p$, if the condition evaluates to *true*, p is executed and otherwise a deadlock occurs. Thus, `[c]p` and $c \rightarrow p$ are trace equivalent. These examples of similar constructs from POOSL and mCRL2 reiterate the importance of formal semantics in transformations between languages.

In the literature, some authors have reported similar experiences about transformations between general-purpose modeling languages. A partial transformation from the hybrid modeling formalism Chi 2.0 [25] to mCRL2 is proposed in [22]. The intention is to enable formal verification on models in Chi 2.0. Semantic differences between Chi 2.0 and mCRL2 makes the transformation complex and hard to maintain. The generated models are also complex and sometimes difficult to analyze by tools. In [12] the authors use a model checker and a theorem prover for verification of UML diagrams. Representing the semantics of UML diagrams in these tools and the consistency of these representations are among the main challenges in this work.

The *Pedal Handling* DSL is defined for a narrow domain and its semantics is less elaborate compared with POOSL, mCRL2, and AML. Thus, it requires less effort to construct the transformations from the DSL to different analysis models.

9 Related Work

In [6] the authors prototype the semantics of a DSL called SLCO. The core semantics is described in terms of a transformation to an intermediate language called CS. Afterwards, CS models are transformed to labeled transition systems and are inspected manually or analyzed by existing tools. During the development of SLCO, the authors have implemented a number of transformations, including transformations from SLCO to SLCO models with equivalent behaviors. The authors suggest that the prototype semantics can be used to compare the underlying labeled transition systems of the source and target models.

The B method [5] has been used in [9] to develop process schedulers based on specifications in a DSL. The information given by a DSL model is taken into account at several refinement steps in B machines. The authors also introduce a decidable logic for expressing proof obligations of the refinement steps. This allows them to automatically prove the refinements.

The mentioned studies focus on validating refinement steps. In contrast, we offer various formal techniques to assess the semantic correctness of different types of artifacts.

In [21] a combination of model-based techniques is used to develop a software bus in a two-phase process. In the first phase, an mCRL2 model of the component is created and validated through simulation. After developing the component, the mCRL2 model is used for model-based testing of the implementation. In the second phase, different properties are verified against the mCRL2 model. The model is improved based on the results and is used for model-

based testing against a second implementation. In comparison, our approach is centered around domain-specific models and artifacts generated from them.

In [13] the POOSL language is used to create a requirements model and a preliminary design model for *Pedal Handling*. These models are translated to CSP [20] and compared using the refinement checker FDR2 [16]. In [13] it is assumed that *Pedal Handling* consists of an input buffer with certain scheduling rules and a component that processes the inputs from the buffer to make decisions about X-ray requests to the tubes. Hence, the scope of the models is wider compared to the work presented in this paper. As a result, refinement checking is not always feasible. In contrast, specifications in the *Pedal Handling* DSL focus on the decision making aspect of the component. This reduces the complexity of the models and makes it feasible to compare behaviors by equivalence checking.

In Section 6, we have used model-based testing to validate an implementation of a DSL model. As an alternative, one can use a model learning approach. Model learning techniques extract an automaton model for an implementation by systematically performing tests on it and observing its behavior [7]. The extracted model can be compared with the labeled transition system described by a DSL model.

10 Conclusions

A DSL allows us to use models that are naturally aligned with the way domain experts reason about a software component. Existing tools enable language designers to define DSLs and to construct transformations to implementation code and analysis models. However, the semantic correctness of the generated artifacts is usually overlooked.

To resolve conflicting interpretations of a DSL, we use a formal semantics of the language. We also propose to have additional mechanisms to validate the correctness of the generated artifacts with respect to the semantics of the DSL.

Introducing redundancy for validating the correctness of the generated artifacts requires extra effort to develop techniques that enable these checks and to analyze the results. However, these redundant mechanisms can effectively detect mistakes in the artifacts. Discovering mistakes through redundant mechanisms raises awareness about the potential inconsistencies between transformations from a DSL and induces search for other forms of redundancy to increase the reliability of the generated artifacts.

We have experimented with this approach as preparation for the redesign of a clinical X-ray generator. We plan to extract models from the existing implementation using model learning techniques. A learned model can give insight in the implemented behavior, and can be compared with the labeled transition system described by a DSL model.

Acknowledgement This research was supported by the Dutch national program COMMIT and carried out as part of the Allegio project.

References

[1] Axini b.v. `http://www.axini.nl`. Visited: August 2015.

[2] mCRL2. `http://mcrl2.org`. Visited: August 2015.

[3] Xtend. `http://www.eclipse.org/xtend/`. Visited: August 2015.

[4] Xtext. `http://www.eclipse.org/Xtext/`. Visited: August 2015.

[5] J.-R. Abrial, J.-R. Abrial, and A. Hoare. *The B-book: assigning programs to meanings*. Cambridge University Press, 2005.

[6] S. Andova, M. van den Brand, and L. Engelen. Prototyping the semantics of a DSL using ASF+SDF: Link to formal verification of DSL models. In *Proceedings of AMMSE'11*, volume 56 of *EPTCS*, pages 65–79, 2011.

[7] D. Angluin. Learning regular sets from queries and counterexamples. *Information and computation*, 75(2):87–106, 1987.

[8] J. Bergstra and J. Klop. Process algebra for synchronous communication. *Information and control*, 60(1):109–137, 1984.

[9] J. Bodeveix, M. Filali, J. Lawall, and G. Muller. Formal methods meet domain specific languages. In *Proceedings of IFM'05*, pages 187–206. Springer, 2005.

[10] H. Ehrig and C. Ermel. Semantical correctness and completeness of model transformations using graph and rule transformation. In *Proceedings of ICGT'08*, volume 5214 of *LNCS*, pages 194–210. Springer-Verlag, 2008.

[11] J. Groote and M. Mousavi. *Modeling and Analysis of Communicating Systems*. MIT-press, 2014.

[12] J. Hooman, H. Kugler, I. Ober, A. Votintseva, and Y. Yushtein. Supporting UML-based development of embedded systems by formal techniques. *Software & Systems Modeling*, 7(2):131–155, 2008.

[13] J. Hooman, A. Mooij, and H. van Wezep. Early fault detection in industry using models at various abstraction levels. In *Proceedings of iFM'12*, pages 268–282. Springer, 2012.

[14] A. Khabbaz Saberi, J. Groote, and S. Keshishzadeh. Analysis of path planning algorithms: a formal verification-based approach. In *Proceedings of ECAL'13*, volume 12, pages 232–239, 2013.

[15] D. Kozen. Results on the propositional μ-calculus. *Theoretical computer science*, 27(3):333–354, 1983.

[16] F. S. E. Ltd. Refinement, failures-divergence - FDR2 user manual. 2010.

[17] A. Mooij, J. Hooman, and R. Albers. Gaining industrial confidence for the introduction of Domain-Specific Languages. In *Proceedings of IEESD'13*, pages 662–667. IEEE, 2013.

[18] G. Plotkin. A structural approach to operational semantics. *Journal of Logic and Algebraic Programming*, 60(61):17–139, 2004.

[19] POOSL. `http://poosl.esi.nl`. Visited: August 2015.

[20] A. Roscoe. *Understanding concurrent systems*. Springer-Verlag New York Inc., 2010.

[21] M. Sijtema, A. Belinfante, M. Stoelinga, and L. Marinelli. Experiences with formal engineering: Model-based specification, implementation and testing of a software bus at Neopost. *Science of computer programming*, 80:188–209, 2014.

[22] F. Stappers. *Bridging Formal Models: An Engineering Perspective*. PhD thesis, Eindhoven University of Techonology, 2012.

[23] B. Theelen, O. Florescu, M. Geilen, J. Huang, P. van der Putten, and J. Voeten. Software/hardware engineering with the parallel object-oriented specification language. In *Proceedings of MEMOCODE'07*, pages 139–148. IEEE, 2007.

[24] J. Tretmans. Model based testing with Labelled Transition Systems. In *Formal methods and testing*, volume 4949 of *LNCS*, pages 1–38. Springer, 2008.

[25] D. van Beek, A. Hofkamp, M. Reniers, J. Rooda, and R. Schiffelers. Syntax and formal semantics of chi 2.0. *Eindhoven University of Technology, Technical Report*, 2008.

[26] L. van Bokhoven. *Constructive Tool Design for Formal Languages*. PhD thesis, Eindhoven University of Techonology, 2002.

[27] M. van den Brand and J. Groote. Software engineering: Redundancy is key. *Science of Computer Programming*, 97:75–81, 2013.

[28] A. van Deursen, P. Klint, and J. Visser. Domain-Specific Languages: an annotated bibliography. *SIGPLAN Notices*, 35(6):26–36, 2000.

[29] J. Verriet, H. Liang, R. Hamberg, and B. van Wijngaarden. Model-driven development of logistic systems using domain-specific tooling. In *Proceedings of CSD&M*, pages 165–176. Springer, 2013.

A Validation Properties for *Pedal Handling*

In this appendix we provide some example properties that can be verified against specifications in the *Pedal Handling* DSL. All the properties presented in this appendix are expressed in a variant of the modal μ-calculus. This logic is the language used by the mCRL2 toolset [2] for property specification. The interested reader can refer to [11] for more details about the syntax and semantics of this logic.

We provide example properties for the DSL model of Fig. 3. As discussed in Section 3, a specification in the Pedal Handling DSL declares a set of input actions and describes the way the state of the system evolves as it perform the actions. For example the DSL model of Fig. 3 declares the input actions *FRFluoOn*, *FRFluoOff*, *StartCond*, and *ResetStartCond*. To express properties for this DSL model, we first explain the behaviors represented by the declared actions:

- *FRFluoOn*: request to start capturing real-time images;
- *FRFluoOff*: request to stop capturing real-time images;
- *StartCond*: notification to prevent X-ray from starting;
- *ResetStartCond*: notification to allow X-ray generation for new requests.

Pedal Handling performs actions $output(xr, p)$ to send X-ray requests for $xr \in XRay$ to $p \in Plane$.

In what follows, we formalize some example properties for the DSL model of Fig. 3. The formalized properties assume that there is a data specification defining the data types *XRay* and *Plane* (see Section 4.1) and refer to the actions discussed above.

- In every reachable state there is at least one outgoing transition (deadlock-freedom):

$$[true^*]\langle true\rangle true$$

- no X-ray is generated from the planes when there is no request from the user:

$$\begin{aligned}
&\nu X(f : \mathbb{B} = false).(\\
&\quad [FRFluoOn]X(true) \wedge\\
&\quad [FRFluoOff]X(false) \wedge\\
&\quad \Big(\forall xr : XRay, p : Plane.[output(xr,p)]\Big(X(f) \wedge\\
&\qquad\qquad (\neg f \Rightarrow (xr = Standby \wedge p = None))\Big)\Big) \wedge\\
&\quad [\overline{FRFluoOn} \wedge \overline{FRFluoOFF} \wedge (\forall xr : XRay, p : Plane.\overline{output(xr,p)})]X(f))
\end{aligned}$$

- if there is a start condition (i.e., a *StartCond* notification that is not canceled by a *ResetStartCond*), X-ray generation cannot be started. It should be noted that a *StartCond* notification does not interrupt the X-ray generation:

$$
\begin{aligned}
&\nu X(fr : \mathbb{B} = false, sc : \mathbb{B} = false).(\\
&\quad [FRFluoOn]X(true, sc) \wedge \\
&\quad [FRFluoOff]X(false, sc) \wedge \\
&\quad [StartCond]X(false, true) \wedge \\
&\quad [ResetStartCond]X(false, false) \wedge \\
&\quad \Big(\forall xr : XRay, p : Plane.[output(xr, p)]\Big(X(false, sc) \wedge \\
&\qquad\qquad \big((fr \wedge sc) \Rightarrow (xr = Standby \ \wedge \ c = None)\big)\Big)\Big)
\end{aligned}
$$